

Index

- 3D rendering, 207
- Academic Computer Center CYFRONET AGH, Kraków, Poland, xxvi
- Academic Computer Center, Gdańsk, Poland, xxvi
- accelerator, 298
- address
 - physical, 184, 310
 - translation lookaside buffer, 184
 - virtual, 183
- address space, 5, 24, 182, 183, 305, 309
- algorithm
 - approximate, 217
 - network model, 92
 - matrix–matrix multiplication, 94, 95
 - matrix–vector multiplication, 93, 236
 - minimum graph bisection, 217
 - prefix computation, 99
 - reduction, 96
 - sorting, 228
- parallel, 63, 123, 125, 316
 - absolute and relative speedup, 64
 - array packing, 119
 - coarse-grained, 142
 - Cole’s sorting, 104
 - communication cost, 64, 85, 183
 - cost, 63, 66, 308
 - efficiency, 63, 66, 157–159, 161, 310
 - fine-grained, 134, 293
 - matrix transpose, 120
 - matrix–matrix multiplication, 85, 288
 - memory requirement, 307
 - minimum graph bisection, 217
 - odd–even transposition sort, 113, 122, 210
 - overhead, 65–67, 70–72, 139, 144
 - parallel running time, 64, 100, 307
 - performance metric, 64, 66, 77, 84, 85, 94, 97, 317
 - portability, 64
 - prefix computation, 80, 98, 118–120, 287
 - prefix computation with segmentation, 119
 - processor complexity, 74, 77, 81, 100, 318
 - reduction, 121
 - scalability, 67, 123, 158–160
 - scaled speedup, 71, 311
 - size of input data, 65
 - speedup, 63, 319
 - tree traversal, 121
- perfect shuffle, 202
- PRAM model, 72
 - array packing, 119
 - finding minimum element, 73
 - finding sum of elements, 75, 76
 - matrix–matrix multiplication, 85, 288
 - prefix computation, 81, 287
 - prefix computation with segmentation, 119
 - sorting, 83, 228
- randomized, 145, 155, 161
- round-robin, 145, 153, 155, 161, 262
- sequential, 1
 - array packing, 119
 - bubble sort, 122
 - counting sort, 83, 103
 - Euclidean, 58, 285
 - Horner’s, 36, 39
 - insertion sort, 113
 - matrix transpose, 119
 - memory requirement, 39, 63, 307
 - merge sort, 104
 - minimum graph bisection, 217
 - performance metric, 35, 38
 - prefix computation, 287
 - prefix computation with segmentation, 119
 - quicksort, 134
 - running time, 38, 39, 63, 307
 - size of input data, 38
 - spatial modeling, 43
- Amdahl’s law, *see* law: Amdahl’s
- American National Aeronautics and Space Administration (NASA), xxi
 - Ames Research Center, xxii
 - Goddard Space Flight Center, 186
- API, *see* application programming interface
- application programming interface, 243, 248, 262, 273, 275, 281
- approximation, 162
- arithmetic-logic unit, 177, 178, 321
- floating-point, 207, 208

344 **Index**

- batch mode, 24
- Bellman Richard Ernest, 170
- Beowulf, *see* cluster: Beowulf
- bijection, 60
- binary tree, *see* interconnection network: binary tree
- bitonic sequence, 113
- bitonic sorting network, *see* network: sorting: bitonic
- blocking, *see* semaphore: blocked-queue or blocked-set
- Brent’s theorem, 79
- broadcasting, *see* message broadcast, and interconnection network: cube: message broadcast
- buffer, 11
 - cyclical, 13
- bus, *see* interconnection network: bus
- busy waiting, 27
- butterfly, *see* interconnection network: butterfly
- characteristic vector, 128
 - base sequence, 128
- checklist, 159–161
- circuit
 - integrated, 207
 - logic, 56, 57, 58, 102, 103
 - VLSI, 61, 196
- cluster, 24, 183, 184, 207, 212, 215, 243, 306
 - Beowulf, 186, 212, 306
 - computer, 185, 306
 - computing, 294, 310
 - constellation, 184, 306, 308
 - data references, 185
 - data transmission
 - latency, 141, 187, 312
 - rate, 187
 - load balancing, 188
 - massively parallel, 313
 - multicore processor, 185, 299, 306
 - node, 184, 306
 - packaging
 - compact, 187
 - slack, 187
 - reliability, 188
 - scalability, 186
 - symmetric multiprocessor, 184, 306
- coherence, *see* memory: cache: data consistency, and OpenMP: data consistency
- collective communication, *see* MPI library: collective communication, and communication between processes: collective
- communication between processes, 2, 157, 158
 - asynchronous, 11, 156
 - blocking, 156, 223
 - buffered, 222
 - collective, 159
 - cost, 159, 160, 294
 - global, 158
 - local, 158
 - message passing, 3, 35, 41, 184, 214, 215, 220, 306, 313
 - message delivery time, 3
 - nonblocking, 222
 - point-to-point, 94, 206, 225
 - shared memory, 3, 35
 - shared variable, 3
 - operation, 3
 - synchronous, 11, 156, 221
- communication channel, 3, 41, 156, 198, 306, 312
 - data transfer rate, 42, 198, 309
 - latency, 94
 - transmission time, 94
- communication complexity, *see* algorithm: parallel: communication cost
- communication properties, *see* interconnection network: communication properties
- company
 - AMD/ATI, 181, 207
 - IBM, 207
 - Inmos, 214, 241
 - Intel, 178, 181, 212
 - Nvidia, 207
 - Sony, 207
 - Sun, 207
 - Toshiba, 207
- comparator, *see* network: comparator
- comparator network, *see* network: comparator
- complex plane, 164
- complexity class, 100
 - L, 118
 - NC, 100, 117, 315
 - subclasses NC^j, 100, 315
 - NL, 118
 - NP, 100
 - NP-complete, 102, 144, 145, 147, 161, 216
 - P, 100
 - P-complete, 65, 100–103, 124, 317
- compression and decompression, 207
- computation
 - asynchronous, 3, 93, 181, 187, 191, 196, 216
 - synchronization operation, 216
 - coarse-grained, 137, 142, 187, 311
 - computer graphics, 43
 - dataflow, 189, 191, 196
 - distributed, 24, 187, 214
 - detecting termination of distributed computation, 154
 - fine-grained, 134, 137, 153, 155, 158, 186, 191, 207, 309, 311
 - flow, 194
 - grain size, 137
 - granularity, 168, 173
 - linear algebra, 43
 - loosely synchronous, 216
 - message-passing, 243
 - multithreaded, 208, 212
 - parallel, 35, 42, 58, 64, 66, 70, 100, 117, 131, 135–137, 156, 189, 214, 309
 - GPUs, 207
 - overhead, 66, 140, 144, 159
 - pipeline, 128, 175, 317
 - pipelining, 125, 128, 317
 - prefix, 80, 118, 120
 - redundant, 140, 142, 158
 - shared-memory, 243
 - synchronous, 40, 94, 96, 180, 216
- computational
 - power, 47
 - step, 36, 38, 40, 64

- computational complexity
 - memory, 38, 39, 63, 307
 - time, 38, 39, 63, 64, 100, 307
- computer
 - Alpha 21364, 62
 - BBN Butterfly, 209
 - Blue Gene/L, xxiv, 43, 62
 - Bull NovaScale, 209
 - C.mmp, 209
 - Compaq ES40, xxiii
 - Compaq GS160, xxiv
 - conventional, 175, 177, 195, 196
 - Convex Exemplar, 209
 - Cosmic Cube, 61
 - CPP DAP Gamma II, 207
 - Cray T3D, 62
 - Cray T3E, 62
 - Cray X1E, xxii
 - Cray XT4, xxii, xxiii
 - Cray XT5, xxiii
 - dataflow, 189, 213, 309
 - control flow, 189, 191
 - Cyberflow, 196
 - data flow, 190
 - DDP, 195
 - EDDN, 196
 - EM-5, 195
 - Manchester, 195, 213
 - MIT, 195, 213
 - Monsoon, 195, 213
 - SIGMA-1, 195
 - structure, 191
 - Earth Simulator, xxii
 - EDVAC, 35
 - Fujitsu VPP 500, 209
 - HP Superdome, 209
 - hybrid, 298
 - IBM Bluefire, xxiii
 - IBM eServer 325, 236
 - IBM p575, xxii
 - IBM p690, xxii
 - IBM RP3, 209
 - Illiac IV, 61, 207
 - Intel DELTA, 62
 - iPSC, 62
 - J-machine, 62
 - K, xxiii
 - Kendall Square Research KSR1, 212
 - MasPar MP-1, 207
 - massively parallel, 313
 - Meiko CS-2, 209
 - MIMD, 182
 - distributed memory, 183, 314
 - loosely coupled, 183
 - shared memory, 182, 209, 314
 - symmetric multiprocessor, xxiii, 180, 182, 184, 200, 320
 - tightly coupled, 182
 - MISD, 198
 - multicomputer, 183
 - multiprocessor, 27, 28, 160, 181, 183, 207, 212, 314
 - distributed memory, 183, 314
 - loosely coupled, 183
 - shared memory, 182, 209, 281, 282, 314
 - tightly coupled, 182
 - nCUBE, 62
 - NEC Cenju-4, 209
 - NYU Ultracomputer, 209
 - personal (PC), 5, 180, 186, 196
 - Pleiades, xxii
 - processor array, 180, 207, 212, 317
 - Quadrics Apemille, 207
 - Roadrunner, 298
 - SGI Altix, xxi, 209
 - SGI Origin 2000, 209
 - SIMD, 181, 207
 - activity mask, 181
 - control processor, 180
 - diversification of computations, 181
 - execution of if-then-else instruction, 181
 - image processing, 181
 - processing element, 180, 317
 - processor array, 317
 - simulation of atmospheric phenomena, 181
 - SISD, 175, 311, 321
 - Solomon, 61
 - Sun Ultra HPC Server, 209
 - Sun V40z, 236
 - supercomputer, 186, 209, 296, 298, 319
 - systolic, 198, 213, 320
 - Warp and iWARP, 213
 - Thinking Machines CM-1, CM-2, CM-5, 207, 209
 - Tianhe-2, xxi, xxv, 209, 296
 - unconventional architecture, 189
 - uniprocessor, 311, 321
 - vector, 177, 321
- computer architecture, 35, 175, 307
 - Cell Broadband Engine, 207
 - conventional, 177
 - dataflow, 189, 212, 309
 - MIMD, 182, 311
 - MISD, 198, 311
 - parallel, 212
 - processor array, 180
 - SIMD, 180, 181, 311
 - SISD, 175, 311
 - systolic, 196, 320
 - unconventional, 189
 - uniprocessor, 212
 - vector, 321
 - von Neumann, 35, 177, 212
- computer game console, 180
 - Sony PlayStation 3, 207
- computer graphics, 207, 209
- computer vision, 207
- computing node, *see* cluster: node
- concurrent processes, 2, 3, 34, 308
 - contention, 7
 - creation
 - dynamic, 216
 - static, 216
 - critical section, 8–10, 12, 18, 27, 271, 308
 - deadlock, 7, 10, 14–16, 34, 309
 - multithreaded, 5, 244
 - mutual exclusion, 7, 10, 271, 315
 - nonterminating, 6
 - parallel execution, 3
 - priority, 5
 - pseudo-parallel execution, 4, 5, 142, 143, 314
 - starvation, 7, 10, 14, 15, 18, 24, 27, 30
 - synchronization, 3, 12, 15, 26, 34, 314
 - condition (event), 13, 17

346 Index

- concurrent program, 1, 3, 24, 34, 307
 - correctness, 6, 34
 - fairness, 7
 - fairness (weak, strong), 310
 - liveness, 6, 7, 10, 15, 34, 312
 - safety, 6, 7, 10, 16, 34, 318
 - condition synchronization, *see* concurrent processes: synchronization: condition (event)
 - condition variable, *see* monitor: condition variable
 - conditional compilation, 248
 - constant of gravitation, 166
 - constellation, *see* cluster: constellation
 - contention
 - resource, 8, 27, 42, 152, 308
 - context switch, *see* operating system: context switch
 - core, 179, 308
 - cost, *see* algorithm: parallel: cost
 - criterion of cost
 - logarithmic, 39
 - uniform, 39
 - critical section, *see* concurrent processes: critical section, and problem: critical section
 - crossbar switch, *see* interconnection network: crossbar switch
 - csh shell, 247
 - CSP notation, 214, 241
 - cube, *see* interconnection network: cube
 - cube connected cycles, *see* interconnection network: cube connected cycles
 - CUDA environment
 - device function (`__device__`), 208
 - general-purpose GPU processing (GPGPU), 209
 - grid, 209
 - host function (`__host__`), 208
 - kernel function (`__global__`), 208
 - thread block, 209
 - thread index, 209
 - cycle, 102
 - Euler, 88
 - data
 - local, 156
 - nonlocal, 156
 - redundant, 158
 - data consistency (coherence), *see* memory: cache: data consistency, and OpenMP: data consistency
 - data dependency, 65, 176, 189, 247
 - data parallelism, *see* designing parallel algorithms: method: data parallelism, and designing parallel algorithms: decomposition: data
 - database, 8, 17, 33
 - transaction, 17
 - dataflow computation model, *see* model of computation: parallel: dataflow
 - de Bruijn Nicolaas Govert, 61
 - deadlock, *see* concurrent processes: deadlock
 - decomposition, *see* designing parallel algorithms: decomposition
 - degree of concurrency
 - average, 126, 130, 139
 - maximum, 126, 130
 - of a problem, 64, 65, 126, 149, 151, 294, 309
 - of a program, 126, 247, 309
 - of an algorithm, 140
 - depth of a network, *see* network: depth
 - depth-first search, 90
 - design checklist, 158
 - designing parallel algorithms, 125
 - assigning tasks to processors, 140, 143, 145, 146, 150, 151, 155, 160, 161
 - block, 146
 - block-cyclic, 146
 - cyclic, 146
 - cost, 168, 173
 - decomposition, 125, 158, 168, 173, 292, 317
 - blockwise, 145
 - data, 126, 129, 131, 147, 158, 161
 - exploratory, 135, 158
 - fine-grained, 158
 - functional, 126, 127, 130, 150, 151, 158, 161
 - mixed (hybrid), 137, 158
 - recursive, 133, 158
 - speculative, 136, 158
 - load balancing, 143, 161, 173, 312
 - centralized, 151, 161
 - decentralized, 152, 161
 - distributed, 152, 161
 - dynamic, 151, 161
 - static, 145, 161
- method, 125
 - data parallelism, 147, 181
 - Foster's, 156
 - functional parallelism, 125–130, 137, 145, 149–151, 154, 158, 161
 - master-worker, 155, 162, 217, 281
 - pipeline, 128, 155
 - producer and consumer, 155
 - task pool, 154
 - task dependency graph, 127–130, 151
- determinism, 58, 103, 118
- differential equation, 166
- directed forest, 120
- distributed program, 3, 24, 310
- distributed shared memory, 183, 212, 310
 - data transfer
 - latency, 184, 312
 - message passing, 184, 310, 313
- distribution
 - data between memories, 183
 - matrix elements, 94
- domain decomposition, *see* designing parallel algorithms: decomposition: data
- effect
 - Amdahl, 65, 123, 305
- efficiency, *see* algorithm: parallel: efficiency
- efficiently parallelizable problem, *see* problem: efficiently parallelizable
- embarrassingly parallel problem, *see* problem: embarrassingly parallel
- embedding, *see* interconnection network: embedding
- encapsulation, 19
- environment OpenCL, 208
- equivalence of computation models, *see* model of computation: parallel: equivalence of models
- Eratosthenes of Cyrene, 126

- Ethernet
 - fast, 187
 - gigabit, 187
- Euclid, 58
- Euler Leonhard, 88, 167
- Euler path, 91, 289
- event synchronization, *see* concurrent processes:
 - synchronization: condition (event)
- fairness, *see* concurrent program: correctness:
 - fairness
- fastest supercomputers, xxi, 181, 209, 296, 298
- fat tree, *see* interconnection network: fat tree
- FIFO queue, 7, 10, 152, 156
- Flynn's taxonomy, 175, 206, 212, 311
 - data stream, 175
 - instruction stream, 175
- fork-join parallelism, *see* OpenMP interface:
 - model of computation: fork-join
 - parallelism
- Foster's method, *see* designing parallel algorithms:
 - method: Foster's
- fractal, 164
- fractal geometry, 164
- function
 - Boolean, 56, 102
 - isoefficiency, 68, 123, 312
 - growth rate, 318
- functional parallelism, *see* designing parallel algorithms: decomposition: functional, and designing parallel algorithms: method: functional parallelism
- functional programming language, *see* programming language: functional
- galaxy, 168
- game theory, 163
- GNU, *see* project: GNU
- granularity, 125, 134, 137, 138, 185, 311
 - coarse, 137, 294
 - fine, 134, 138, 185, 293
 - grain size, 142, 311
 - medium, 137
- graph
 - directed, 88, 120
 - Eulerian, 88
 - loop, 120
 - vertex adjacency lists, 89
- graph bisection, *see* problem: minimum graph
 - bisection
- Gray code, 54
- greatest common divisor, *see* algorithm:
 - sequential: Euclidean
- hazard
 - control, 177
 - data, 176
- Heron of Alexandria, 189
- Horner William George, 36
- hypercube, *see* interconnection network: cube
- I/O
 - operation, 4, 24, 36, 142, 269
 - port, 156
- image processing, 129, 207
 - filter
 - averaging, 130
 - Gauss, 130
 - lp2, 130
 - filtering, 129
 - pixel, 129, 164, 300, 303
 - resolution, 165
 - smoothing, 130
- induction, 54, 116
- Infiniband, *see* interconnection network: switch:
 - Infiniband
- injection, 50, 52, 54
- instruction
 - atomic, 27
 - compare-and-swap, 27
 - empty, 27
 - exchange, 27
 - fetch-and-add, 27
 - test-and-set, 26
 - vector, 177
- instruction list, *see* processor: instruction list
- instruction pipelining, 317
 - conditional branch, 176
 - control hazard, 177
 - data dependency, 176
 - data hazard, 176
- integration
 - method
 - rectangles, 162, 241
 - trapezoids, 163, 164, 241, 281
 - step, 162
- interconnection network, 41, 180, 182, 312
 - binary tree, 59, 60
 - double-rooted, 60
 - dynamic, 204, 206, 320
 - static, 204, 206, 320
 - bisection bandwidth, 42, 305
 - bisection width, 42, 48, 206, 216, 305
 - bus, 182, 199, 305
 - access algorithm, 199
 - bandwidth, 199, 200
 - cost, 199
 - data transmission latency, 199, 312
 - scaling, 200
 - butterfly, 46, 48, 203, 206, 209, 306
 - communication properties, 42, 43, 53, 59
 - completely-connected, 42, 48, 216
 - contention
 - communication resource, 42
 - cost, 42, 48, 206
 - crossbar switch, 182, 184, 200, 206, 209, 308
 - cost, 200
 - scaling, 200
 - cube, 44, 47, 48, 53, 59, 60, 96, 308
 - k*-dimensional, 44, 48, 53, 55, 60
 - data transfer statement, 97
 - four-dimensional, 285
 - message broadcast, 49, 98, 121
 - three-dimensional, 54, 55
 - cube connected cycles, 46, 60, 309
 - de Bruijn, 61
 - degree, 42, 48
 - diameter, 42, 48, 60, 205, 309
 - dynamic, 199
 - evaluation parameters, 205, 206
 - latency, 206, 312
 - edge connectivity, 42, 48, 206

348 Index

- interconnection network (*cont.*)
 - embedding, 50, 310
 - congestion, 51, 52, 53
 - dilation, 51, 52, 53
 - expansion, 52, 53
 - load factor, 51, 52, 53
 - fat tree, 205, 299, 311, 320
 - local area network, 187
 - maximum degree vertex, 206
 - mesh, 48, 313
 - k*-dimensional, 43, 320
 - multidimensional, 53–55
 - one-dimensional, 43, 52, 53, 204, 320
 - three-dimensional, 43, 60
 - two-dimensional, 43, 52–54, 55, 60, 285
 - mesh of trees, 313
 - two-dimensional, 43, 48
 - message routing, 41, 199, 318
 - routing procedure, 41, 206, 318
 - multistage, 182, 201, 213, 314, 315
 - node, 198, 312
 - omega, 182, 201, 206, 209, 213, 315
 - cost, 203
 - data transmission time, 203
 - message routing, 202, 318
 - proprietary, 187
 - resistance to damage, 42
 - scaling, 45, 199
 - sparse, 42, 43
 - star, 185, 204, 320
 - static, 199
 - evaluation parameters, 42
 - latency, 94
 - switch, 198, 306, 312, 319
 - degree, 198
 - Infiniband, 299
 - message broadcast, 199
 - message buffering, 199
 - ports, 199
 - topology, 41, 47, 180, 199, 320
 - torus, 48, 59, 286, 320
 - k*-dimensional, 320
 - data transfer statement, 94
 - doubly twisted, 59
 - message broadcast, 121, 288
 - one-dimensional, 43, 53, 54, 60, 93, 138, 185, 318
 - three-dimensional, 43
 - two-dimensional, 44, 59, 94, 286
 - tree, 204, 320
 - vertex connectivity, 42, 206
 - wide area network, 187
 - Interdisciplinary Centre for Mathematical and Computational Modelling, University of Warsaw, Poland, xxvi, 236
 - interleaving, *see* operating system: interleaving
 - isoefficiency function, *see* function: isoefficiency
 - Karp–Flatt metric, *see* sequential fraction of parallel computation
 - L, *see* complexity class: L
 - LAN, *see* interconnection network: local area network
 - latency, 3, 94, 142, 206, 207, 312
 - law
 - Amdahl’s, 69, 123, 305
 - Gustafson–Barsis’s, 70, 123, 311
 - Moore’s, 178, 212, 314
 - level of abstraction, 37, 38, 42, 58
 - library
 - BLAS, 208
 - DirectX, 208
 - java.util.concurrent, 9
 - MPI, 214
 - Pthreads, 9, 244, 282
 - PVM, 187, 214, 241
 - linear algebra, 43
 - link, *see* communication channel
 - list, 87, 120
 - Green 500, 298
 - Top 500, 181, 209, 296
 - list structure, 87
 - liveness, *see* concurrent program: correctness:
 - liveness
 - load balancing, *see* designing parallel algorithms:
 - load balancing
 - local area network, *see* interconnection network:
 - local area network
 - locality of data reference, 159, 312
 - spatial, 178
 - temporal, 178
 - logarithmic cost criterion, *see* criterion of cost:
 - logarithmic
 - logic circuit, *see* model of computation: parallel:
 - logic circuit
 - Mandelbrot Benoît B., 164
 - massively parallel processing system, *see* cluster:
 - massively parallel
 - memory
 - access
 - random, 36
 - sequential, 36
 - cache, 27, 177, 179, 184, 200, 212, 273, 282, 306, 310
 - ccNUMA, 185, 209
 - data consistency, 185, 249, 273, 282, 306
 - hit, 178, 306
 - line, 178, 306, 312
 - miss, 178, 306
 - spatial locality, 199, 312
 - temporal locality, 178, 312
 - distributed, 34, 183, 309
 - DRAM, 178
 - global, 207
 - access time, 207
 - hierarchical, 141, 177, 184, 311
 - local, 39, 93, 156, 181, 183, 184
 - nonlocal, 142, 184
 - nonuniform memory access, xxiii, 141, 182, 184, 315
 - shared, 3, 6, 34, 39, 183–186, 207, 243, 244, 319
 - access time, 207
 - bandwidth, 182, 313
 - module, 182, 199
 - uniform memory access, 182, 184, 321
- memory complexity, *see* algorithm: sequential:
 - memory requirement, and algorithm:
 - parallel: memory requirement

- memory requirement, *see* algorithm: sequential:
 - memory requirement, and algorithm:
 - parallel: memory requirement
- memory wall, *see* problem: memory wall
- mesh, *see* interconnection network: mesh
- message broadcast, 241
 - all-to-all, 121, 288
 - one-to-all, 98, 121
- message passing, *see* communication between
 - processes: message passing
- method
 - divide and conquer, 170
 - dynamic programming, 170
 - common subproblems property, 170
 - optimal substructure property, 170
 - Eulerian cycle, 88
 - Monte Carlo, 162, 241, 281, 304
 - convergence, 163
 - Newton's, 194
 - pointer jumping, 87, 120
- Milky Way-2, *see* computer: Tianhe-2
- MIMD (MISD), *see* computer architecture:
 - MIMD (MISD), and computer: MIMD (MISD)
- model of computation, 35
 - parallel
 - arbitrary CRCW PRAM, 48
 - BSP, 56
 - combining CRCW PRAM, 40, 120, 288
 - common CRCW PRAM, 40, 48, 81
 - comparator network, 112, 307
 - comparison of network models, 50
 - comparison of PRAM models, 47
 - conflicts in shared memory access, 40
 - CRCW PRAM, 40, 47, 82
 - CREW PRAM, 40, 47, 58, 83, 85, 228
 - dataflow, 189
 - equivalence of models, 58
 - EREW PRAM, 40, 47, 74
 - LogGP, 56
 - logic circuit, 56, 58, 102, 313
 - LogP, 56
 - network model, 35, 41, 187, 315
 - PRAM, 35, 39, 316
 - priority CRCW PRAM, 48
 - processor number, 40
 - shared memory, 316
 - sorting network, 56, 307
 - sequential RAM, 35, 58, 189, 318
 - computational step, 36
 - random access memory, 36
- module, 19
 - memory, 182, 199
 - TriBlade, 299
- monitor, 18, 25, 314
 - condition variable, 19, 314
 - operations wait and signal, 19
- Moore Gordon E., 178
- MPI library, 79, 187, 224, 241, 243, 282, 299
 - MPI_Aint, 302
 - MPI_Allgather, 231
 - MPI_Allreduce, 227
 - MPI_ANY_SOURCE, 223
 - MPI_ANY_TAG, 223
 - MPI_2INT, 227
 - MPI_BAND, 227
 - MPI_Barrier, 236
 - MPI_Bcast, 224, 225, 241
 - MPI_BOR, 227
 - MPI_BXOR, 227
 - MPI_BYTE, 221
 - MPI_CHAR, 221
 - MPI_COMM_WORLD, 220, 228, 230
 - MPI_Comm, 221, 222, 225, 228, 230, 231, 234
 - MPI_Comm_free, 230
 - MPI_Comm_rank, 220, 224
 - MPI_Comm_size, 220, 224
 - MPI_Comm_split, 228, 230
 - MPI_Datatype, 221–223, 225, 230, 231, 234
 - MPI_DOUBLE, 221
 - MPI_DOUBLE_INT, 227
 - MPI_ERROR, 222
 - MPI_Exscan, 80
 - MPI_Finalize, 219, 224
 - MPI_FLOAT, 221, 300
 - MPI_FLOAT_INT, 227
 - MPI_Gather, 230, 231, 233, 234, 300
 - MPI_Gatherv, 234–236
 - MPI_Get_address, 302
 - MPI_Get_count, 223
 - MPI_Init, 219, 224
 - MPI_INT, 221, 300
 - MPI_Irecv, 143, 223
 - MPI_Isend, 143, 222, 223
 - MPI_LAND, 227
 - MPI_LONG, 221
 - MPI_LONG_DOUBLE, 221
 - MPI_LONG_DOUBLE_INT, 227
 - MPI_LONG_INT, 227
 - MPI_LOR, 227
 - MPI_LXOR, 227
 - MPI_MAX, 227
 - MPI_MAXLOC, 227
 - MPI_MIN, 227
 - MPI_MINLOC, 227
 - MPI_Op, 225
 - MPI_PACKED, 221
 - MPI_PROD, 227
 - MPI_Recv, 220, 222–225, 300
 - MPI_Reduce, 79, 224, 225, 227, 229, 230, 241
 - MPI_Scan, 80
 - MPI_Scatter, 231
 - MPI_Send, 220–222, 224, 225, 300
 - MPI_SHORT, 221
 - MPI_SHORT_INT, 227
 - MPI_SOURCE, 222
 - MPI_Status, 222, 223
 - MPI_SUCCESS, 220
 - MPI_SUM, 227
 - MPI_TAG, 222
 - MPI_Test, 143, 222, 223
 - MPI_Type_commit, 303
 - MPI_Type_create_struct, 302
 - MPI_Wait, 143, 222, 223
 - MPI_Wtime, 235, 236, 239
 - MPI_UNDEFINED, 228
 - MPI_UNSIGNED, 221
 - MPI_UNSIGNED_CHAR, 221
 - MPI_UNSIGNED_LONG, 221
 - MPI_UNSIGNED_SHORT, 221
 - collective communication, 224, 225, 300
 - communicator, 220, 225, 228, 230, 231, 236

350 **Index**

- MPI library (*cont.*)
 - derived datatype, 221, 300
 - handle, 301
 - map, 301
 - signature, 301
- history, 215
- message broadcast, 225
- model of computation, 215
- MPI-1, 215, 241
- MPI-2, 215, 241
- MPI-3, 215
- program compilation and execution, 218
 - mpicc command, 218
 - mpirun command, 218
- MPP, *see* cluster: massively parallel
- multiprocessor, *see* computer: multiprocessor
- multiprogramming, *see* operating system:
 - multiprogramming
- multitasking, *see* operating system: multitasking
- multithreading, 6, 34, 180, 207, 208, 282
- mutual exclusion, *see* concurrent processes:
 - mutual exclusion
- NC, *see* complexity class: NC
- network
 - Beneš, 210
 - comparator, 112, 307
 - de Bruijn, 61
 - depth, 113
 - merging, 115
 - permutation, 210
 - size, 113
 - sorting, 113, 307
 - Batcher's, 113, 116
 - bitonic, 113
- network model, *see* model of computation:
 - parallel: network
- Newton Isaac, 166, 194
- NL, *see* complexity class: NL
- nondeterminism, 118, 144, 192, 196, 249, 251, 253, 257
- nonuniform memory access, *see* memory:
 - nonuniform memory access
- NP, *see* complexity class: NP
- NP-complete, *see* complexity class: NP-complete
- NUMA, *see* memory: nonuniform memory access
- number
 - complex, 164
 - composite, 126, 233
 - prime, 126
 - random, 162, 163, 263, 304
- NYU Ultracomputer, *see* computer: NYU Ultracomputer
- Occam's (or Ockham's) razor, 214
- Ockham, *see* William of Ockham
- OpenMP interface, 208
 - _OPENMP macro, 248
 - Architecture Review Board, 243
 - clause, 246, 251, 255, 257, 259
 - collapse, 255
 - copyin, 251, 268
 - copyprivate, 257, 268
 - default, 251, 259, 261
 - final, 259
 - firstprivate, 251, 255, 257, 259, 260
 - if, 251, 259, 266
 - lastprivate, 253, 255, 257, 261
 - mergeable, 259
 - nowait, 255, 257, 262, 269
 - num_threads, 247, 251, 267
 - ordered, 255
 - private, 246, 251, 255, 257, 259, 260
 - reduction, 251, 255, 257, 264
 - schedule, 255, 262
 - shared, 246, 251, 259, 260
 - untied, 259
 - cOMPunity organization, 243, 281
- construct, 249
 - atomic, 272
 - barrier, 260, 270
 - critical, 260, 265, 271, 275, 276, 278, 304
 - flush, 273
 - master, 269
 - ordered, 271
 - parallel, 250, 254, 267
 - sections, 252, 255
 - single, 252, 257, 268
 - taskwait, 270
 - task, 258, 259
 - loop, 252, 254, 255, 275, 304
 - worksharing, 252
- critical section, 271, 272, 276, 304
- data consistency, 245
- directive, 243
 - threadprivate, 274
- environment variable, 243, 247
 - OMP_NUM_THREADS, 247, 252, 267
- ICVs, 247
- library function, 243, 247, 248, 250, 251
 - omp_set_num_threads, 247
 - omp_get_num_threads, 247
 - omp_get_thread_num, 247, 248, 251
 - omp_get_wtick, 278
 - omp_get_wtime, 278
 - omp_set_num_threads, 252, 267
- memory
 - cache, 245, 249
 - data consistency, 249
 - relaxed-consistency model, 244
 - shared, 244
 - thread's temporary view of memory, 244
 - threadprivate, 245
- model of computation, 244
 - fork-join parallelism, 246
- OpenMP versions 2.5, 3.0, 3.1, 4.0, 243
- program compilation, 252
 - command pgcc, 252
- race condition, 249, 260
- region, 250
 - active, 245
 - inactive, 245
 - parallel, 245, 246, 247, 250–252, 257, 258, 266, 267, 270, 274
- runtime system, 248, 258, 259, 263, 264, 270
- structured block, 250, 251
- synchronization barrier, 251, 252, 254, 257, 260, 268–270
- task, 244, 256, 257
 - child task, 258, 270
 - explicit, 245, 250
 - implicit, 245, 246, 250, 256, 258

- region, 245, 246, 250
- scheduling point, 258, 270
- tied, 245
- team of threads, 247, 250
- thread, 244
 - initial, 245
 - master, 245–247, 250, 251, 258, 269, 274
 - number, 247, 248
 - synchronization, 249, 251, 252, 254, 270
 - unbalanced workloads, 257, 263
- variable
 - nthreads-var*, 247
 - private, 243, 244, 246, 253
 - shared, 243, 246, 249, 253, 260, 270, 273
 - threadprivate, 274
- operand, 189
- operating system, 1, 4, 5, 34, 160, 236, 244, 247, 315
 - concurrency of processes, 4, 34
 - context switch, 4, 207
 - interleaving, 2, 4, 8, 9, 25
 - interrupt, 5
 - priority, 5
 - Linux, 34, 186, 299
 - load balancing, 5
 - multiprogramming, 24, 314
 - multitasking, 314
 - process, 5, 34
 - Solaris-2, 34
 - task scheduling, 5
 - shortest job next principle, 32
 - thread, 5, 34, 244
 - time-sharing, 4, 5, 142, 244
 - Unix, 34
 - Windows, 34
- operation
 - arithmetic, 38
 - associative, 41, 79, 118
 - atomic, 3, 9, 25, 26, 272
 - binary, 79, 118
 - combining, 50
 - communication, 159
 - commutative, 80, 80
 - compare-and-swap, 3
 - comparison, 38
 - computational, 48, 79, 159
 - concatenation, 80
 - control, 36
 - copy, 119
 - data transfer, 196
 - dominant, 38
 - exchange, 3
 - fetch-and-add, 3
 - fixed-point, 144
 - floating-point, xxi, 186, 235, 293, 296, 299, 319
 - gate, 192
 - latency, 142
 - logical (Boolean), 36, 38, 40
 - merge, 192, 192
 - nondeterministic, 192
 - product (Boolean), 80
 - relational, 192
 - SAXPY, 208
 - scan, 80
 - select, 192
 - sink, 192
 - square, 194
 - sum (Boolean), 80
 - swap, 38
 - switch, 193
 - test-and-set, 3
 - vector, 207, 299
- optimal binary search tree, 168, 169–171, 172, 173, 304
- order
 - postorder, 121
 - preorder, 121, 288
 - topological, 102
- overlapping communication and computation, 52, 142, 222, 315
- P, *see* complexity class: P
- P-complete, *see* complexity class: P-complete
- package, 19
- parallel
 - computation thesis, 117, 118
 - execution
 - instruction, 175
 - loop, 260
 - operations, 191, 207
 - tasks, 136, 160
 - programming model, 156
 - Foster's, 156
 - slackness, 143, 208, 317
- parallel execution of processes, *see* concurrent processes: parallel execution
- parallel overhead, *see* algorithm: parallel: overhead
- parallel program, 3, 24, 307
 - correctness, 7, 249
 - fine-grained, 185, 186
 - scalability, 215
 - SPMD method, 216, 242, 248
- parallel running time, *see* algorithm: parallel: parallel running time
- parallel slackness, *see* parallel: slackness
- parallel time complexity, *see* algorithm: parallel: parallel running time
- parallel work, *see* algorithm: parallel: cost
- performance metric, *see* algorithm: parallel: performance metric, and algorithm: sequential: performance metric
- pipelined execution of instructions, 175, 212, 317
- pointer, 87
 - stack, 5
- pointer jumping, 87
- polylogarithmic complexity, 58, 100–102
- portability
 - MPI library, 215
 - parallel algorithm, 64, 159
 - software, 186
- POSIX interface, 244, 282
- Poznań Supercomputing and Networking Center, Poland, xxvi
- PRAM, *see* model of computation: parallel: PRAM
- PRAM-on-chip, 61
- prefix, 80, 98, 118, 120, 287
- prefix computation, *see* problem: prefix computation, and algorithm: sequential: prefix computation, and algorithm: parallel: prefix computation
- preprocessor, 248

352 **Index**

- principle
 - Occam’s razor, 214
 - optimality, 170
 - shortest job next, 32
 - zero-one, 114–116, 123
- problem
 - n*-body, 165
 - MPI program, 241
 - OpenMP program, 281
 - parallel program, 292
 - sequential program, 167
 - NC \neq P?, 100, 124, 315
 - P \neq NP?, 100
 - array packing, 119
 - bin packing, 144
 - cigarette smokers, 33
 - circuit value, 102, 124
 - unrestricted, 124
 - computing approximation of π , 162, 210
 - MPI program, 241
 - OpenMP program, 281, 304
 - sequential program, 162
 - computing integral value, 161
 - MPI program, 240
 - OpenMP program, 281
 - sequential program, 162
 - computing Mandelbrot set, 164
 - MPI program, 241, 299
 - OpenMP program, 281
 - sequential program, 165
 - computing number of descendants, 91
 - computing value of polynomial
 - Horner’s algorithm, 36
 - RAM program, 36
 - constructing optimal binary search tree, 168
 - OpenMP program, 304
 - parallel program, 168
 - constructing optimal binary search tree
 - sequential program, 171
 - critical section, 8, 17, 26, 27
 - decision, 101, 124
 - detecting termination of distributed
 - computation, 154, 173
 - dining philosophers, 14
 - ecoregion map, 146, 173
 - efficiently parallelizable, 58, 100, 173, 315
 - embarrassingly parallel, 126, 173
 - evaluation of multiple-choice test results, 126
 - finding minimum element, 49, 61, 73, 100
 - finding prime numbers, 242, 277
 - MPI program, 232
 - network model, 128
 - OpenMP program, 277
 - pipelining method, 128
 - PRAM, 128
 - sieve of Eratosthenes, 126, 233, 277
 - speedup, 127
 - finding roots of trees, 120
 - functional, 101, 124
 - inherently sequential, 100
 - Königsberg bridges, 88
 - knapsack, 173
 - leader election, 154, 174
 - list ranking, 87
 - load balancing, 5
 - matrix transpose, 119, 121
 - two-dimensional mesh, 121
 - matrix–matrix multiplication, 94, 196
 - combining CRCW PRAM, 120, 288
 - CREW PRAM, 85, 120, 288
 - cube, 288
 - matrix–vector multiplication, 92, 196, 236
 - MPI program, 236
 - maximum cut of graph, 242
 - memory wall, 177, 212
 - minimum cut, 216
 - minimum graph bisection, 216, 275
 - MPI program, 217
 - OpenMP program, 275
 - NC, 100, 124, 315
 - NP, 100
 - NP-complete, 161, 216
 - NP-hard, 163
 - one-lane bridge, 32
 - P, 100
 - P-complete, 100–103, 124, 317
 - prefix computation, 98, 118, 120, 287
 - segmentation, 118
 - sums, 123
 - producer and consumer, 11
 - readers and writers, 17, 30
 - reduction, 79, 123, 225, 264
 - one-dimensional mesh, 121
 - two-dimensional mesh, 121
 - satisfiability of Boolean expressions, 102
 - scheduling, 173
 - shortest schedule, 144
 - size, 38, 64, 158, 160
 - sleeping barber, 33, 284
 - sorting, 83, 100, 122, 174, 228, 276
 - MPI program, 228
 - OpenMP program, 276
 - transforming unrooted tree into rooted tree, 90
 - traveling salesman, 101
- process
 - multithreaded, 282
 - sequential, 1, 5, 319
- processing
 - element, 180, 181, 198, 317
 - unit, 184
- processor
 - AMD Opteron, 236
 - AMD Opteron 2210, 298
 - array, 180
 - ATI Radeon HD 4000/5000, 207
 - Cell, 207
 - dual-core, 179
 - extensions
 - MMX, 181
 - SSE, 181
 - Fermi, 207
 - GPU, 208
 - graphics processing unit, 181, 207
 - IBM POWER6, xxiii
 - IMB PowerXCell 8i, 298
 - instruction list, 26, 37, 181
 - Intel 80-core, 212
 - Kepler, 207
 - Maxwell, 207
 - multicore, 6, 178, 185, 207, 294, 314
 - Pentium, 175

- program counter, 5, 6, 36, 181, 244
- real, 2
- scalar, 177
- streaming, 207
- superscalar, 175
- Tesla, 207
- UltraSPARC T2, 207
- vector, 177, 212, 321
 - pipelined arithmetic-logic unit, 177
- virtual, 1, 2
- processor complexity, *see* algorithm: parallel:
 - processor complexity
- program counter, *see* processor: program counter
- program parallelization
 - automatic, 247, 282
 - incremental, 248
- programming
 - concurrent, 8, 34
 - distributed, 34
 - dynamic, 170
 - multithreaded, 34, 244, 282
 - parallel, 34
 - message-passing, 214
 - shared-memory, 243
- programming language
 - Ada, 10, 25, 34
 - C, 208, 214, 215, 219, 221, 226, 227, 243, 246, 252
 - C++, 208, 215, 243, 246
 - C#, 25
 - Concurrent Euclid, 25
 - Concurrent Pascal, 25
 - Fortran, 214, 243, 246
 - 90, 215
 - 2008, 215
 - HPF, 244, 282
 - functional, 195
 - HASAL, 212
 - Id, 212
 - Java, 25, 34
 - Lapse, 212
 - Mesa, 25
 - Modula 3, 25
 - occam, 214, 241
 - SISAL 1.2, 195, 212
 - VAL, 212
- project
 - GNU, 186
 - Open MPI, 215, 241, 299
- protocol
 - post-protocol, 11, 27, 30
 - pre-protocol, 11, 27, 30
- pseudo-parallel execution of processes, *see*
 - concurrent processes: pseudo-parallel execution
- pseudo-parallel execution of tasks, *see* concurrent processes: pseudo-parallel execution
- Pthreads, *see* library: Pthreads
- PVM, *see* library: PVM
- race condition, *see* OpenMP interface: race condition
- RAM, *see* model of computation: sequential RAM
- RAM program, *see* sequential program: RAM
- random number generator, 304
- random sample, 162
- rank, 104
 - cross, 104
- recursion, 80, 114, 116, 133, 134
- recursive doubling, 80
- reduction, 101
 - complexity of reduction, 101
 - NC, 101, 102
- register, 5, 244
 - arithmetic, 4, 36, 40, 178, 181
 - control, 36
 - program counter, 4
- reliability, *see* cluster: reliability
- replication
 - computation, 142, 159, 160
 - data, 141, 159, 160
- resistance to damage, 24
- resource, 5, 14, 244
- ring, *see* interconnection network: torus:
 - one-dimensional
- routing, 41, 42, 199, 202, 204, 206, 215, 318
- rule
 - firing, 192
 - thumb, 160
- running time, *see* algorithm: sequential: running time
- safety, *see* concurrent program: correctness: safety
- scalability, 67, 123, 215, 318
- scalability of parallel algorithm, *see* algorithm:
 - parallel: scalability
- scheduling, *see* problem: scheduling
- semantic, 249
- semaphore, 8, 24, 128, 318
 - binary, 15, 25
 - blocked-queue, 9, 25
 - blocked-set, 24
 - busy-wait, 24
 - general, 9, 13, 15
 - operations wait and signal, 9
 - simulation of general semaphore, 30, 283
 - split, 14, 30
 - strong, weak, 25
 - waiting queue, 9
- sequential fraction of parallel computation, 71, 123, 312
- sequential program, 160
 - concurrency
 - halting problem, 6
 - correctness, 6, 34
 - assertion, 6
 - partial, 6
 - postcondition, 6
 - precondition, 6
 - total, 6
 - RAM, 36, 285
- serial process, *see* process: sequential
- server, 180
 - computing, 5, 188, 306
 - database, 188, 306
 - www, 188, 306
- set
 - Mandelbrot, 164, 241, 281, 299
 - set image, 164
- sieve of Eratosthenes, *see* problem: finding prime numbers: sieve of Eratosthenes

354 Index

- SIMD (SISD), *see* computer architecture: SIMD (SISD), and computer: SIMD (SISD)
- SIMT, 207
- simulation
- atmospheric phenomena, 181
 - communication operation, 51
 - execution of algorithm, 51, 53, 66, 79
 - general semaphore, 30, 283
 - motion of bodies, 165
 - PRAM computation, 48, 49, 59
 - processor computation, 53
 - simulation step, 165
 - simultaneous reads and writes, 48, 58
 - weather phenomena, 149
- size of a network, *see* network: size
- SMP, *see* computer: MIMD: symmetric multiprocessor, and cluster: symmetric multiprocessor
- software, 186
- free of charge, 186, 215
 - open source, 186
- sorting, *see* problem: sorting, and algorithm: parallel: sorting
- sorting network, *see* network: sorting
- speedup, *see* algorithm: parallel: speedup
- spinning, 27
- SPMD, *see* parallel program: SPMD
- stack, 5, 6, 244
- starvation, *see* concurrent processes: starvation
- statement
- atomic, 272
 - data transfer, 94, 97, 121
 - parfor, 73
 - receive, 93, 121
 - send, 93, 121
- supercomputer, *see* computer: supercomputer
- surjection, 54
- switch, *see* interconnection network: switch
- symmetric multiprocessor, *see* computer: MIMD: symmetric multiprocessor, and cluster: symmetric multiprocessor
- synchronization, *see* concurrent processes: synchronization
- synchronization barrier, 27
- centralized, 28
 - dissemination, 29
 - symmetrical, 29
 - two-task, 28
- system
- computer, 184, 306
 - concurrent, 8
 - distributed, 24, 34
 - parallel
 - scalability, 123
- task, 1, 125, 156, 168, 244, 320
- agglomeration, 157–160
 - granularity, 154, 156
 - coarse, 152
 - fine, 156
 - medium, 137
 - size, 145, 152, 157–159
- task dependency graph, *see* designing parallel algorithms: task dependency graph
- thread, 5, 207, 244, 320
- context switching, 207
 - queue, 207
 - states, 207
- throughput, *see* memory: shared: bandwidth
- time complexity, *see* algorithm: sequential: running time
- time-sharing, *see* operating system: time-sharing
- TLB, *see* address: translation lookaside buffer
- torus, *see* interconnection network: torus
- transputer, 214
- tree
- binary search, 168
 - rooted, 90, 120
 - traversal, 121
 - unrooted, 88
- Turing Alan Mathison, 36, 58
- Turing machine, 36, 58, 103
- type
- derived, 221, 301, 302
 - semaphore, 10
- UMA, *see* memory: uniform memory access
- unary notation, 58
- uniform cost criterion, *see* criterion of cost: uniform
- uniform memory access, *see* memory: uniform memory access
- variable
- local, 74, 76, 81–83, 85, 93, 95, 97, 99, 105
- vector instruction, 177
- video game, 207
- virtual ring, 294
- von Neumann John, 35, 177, 212
- Wallis John, 210
- WAN, *see* interconnection network: wide area network
- wide area network, *see* interconnection network: wide area network
- William of Ockham, 214
- worst-case time complexity, *see* algorithm: parallel: worst-case time complexity
- Wrocław Center for Networking and Supercomputing, Poland, xxvi
- zero-one principle, *see* principle: zero-one